

## NVM Express Technical Errata

<b>Errata ID</b>	031
<b>Change Date</b>	7/18/2012
<b>Affected Spec Ver.</b>	NVM Express 1.0c
<b>Corrected Spec Ver.</b>	

### Submission info

Name	Company	Date
Matthew Wilcox	Intel	7/11/2012
Ken Okin	Virident	7/11/2012
Al Rickey	Intel	7/11/2012
Amber Huffman	Intel	7/11/2012

The natural alignment of a PRP List is 64-bits (Qword). A clarification is added that a PRP List shall be Qword aligned.

A recommendation is added that namespace identifiers remain static to avoid issues with fast boot operations in UEFI or OS drivers.

The UEFI informative appendix is modified to remove material that is more properly referenced to material within UEFI. Refer to NVMe related proposals that are submitted to the UEFI Forum.

Description of the specification technical flaw:

***Modify the paragraph following Figure 11 as shown below:***

The first PRP entry contained within the command may have a non-zero offset within the memory page. The first PRP List entry (i.e. the first pointer to a memory page containing additional PRP entries) that if present is contained in the PRP Entry 2 location within the command, **shall be Qword aligned and** may also have a non-zero offset within the memory page. All other PRP and PRP List entries shall have a memory page offset of 0h, i.e. the entries are memory page aligned based on the value in CC.MPS. The last entry within a memory page, as indicated by the memory page size in the CC.MPS field, shall be a PRP List pointer if there is more than a single memory page of data to be transferred.

***Modify section 6.1 as shown below:***

A namespace is a collection of logical blocks that range from 0 to the capacity of the namespace – 1. The number of namespaces present is reported in the Identify Controller data structure. The namespaces are allocated in order (starting with 1) and packed sequentially. Namespaces identifiers may change across power off conditions. **However, it is recommended that namespace identifiers remain static in order to avoid issues with EFI or OS drivers fast discovery processes.** The management (creation, deletion) of namespaces is outside the scope of this specification.

The Namespace Size field in the Identify Namespace data structure defines the total size of the namespace in logical blocks (LBA 0 through  $n-1$ ). The Namespace Utilization field in the Identify Namespace data structure defines the number of logical blocks currently allocated in the namespace. The Namespace Capacity field in the Identify data structure defines the maximum number of logical blocks that may be allocated at one time as part of the namespace in a thin provisioning usage model. The following relationship holds: Namespace Size  $\geq$  Namespace Capacity  $\geq$  Namespace Utilization.

A namespace may or may not have a relationship to a Submission Queue; this relationship is determined by the host software implementation. The controller shall support access to any valid namespace from any I/O Submission Queue.

***Remove section 10 as shown below:***

***10.1 EFI Information***

~~For additional information on recommendations provided in this appendix, please refer to the following specifications:~~

- ~~• PCI Firmware 3.0 specification (<http://www.pcisig.com>)~~
- ~~• UEFI 2.1 specification (<http://www.uefi.org>)~~

***10.1.1 EFI GUID***

~~The NVM Express revision 1.0 GUID shall be 51116915-C34B-4D8E-86DB-6A70F2E60DAA.~~

***10.1.2 Version Information***

~~EFI modules report their version information as a 16-bit value. The upper byte shall contain the Major Revision number and the lower byte shall contain the Minor Revision number. As an example, 0105h corresponds to a revision number of 1.05.~~

~~EFI Modules shall register their module in the EFI Configuration Table with the GUID listed in section 0. The associated data structure with the GUID shall contain the version information as the first 16 bits.~~

### **10.1.3 EFI Module Discovery**

~~EFI module discovery should be done using provided OS API calls to retrieve the EFI Configuration Table. The specific OS API calls are beyond the scope of this document.~~  
~~Host software should search the EFI Configuration Table for the NVM Express GUID. If a match is found then an NVM Express EFI module was or is currently loaded. The associated data structure may be read to obtain the version information.~~

*Update section 1.10 as shown below:*

## **1.10 References**

PCI specification, revision 3.0. Available from <http://www.pcisig.com>.

PCI Express specification, revision 2.1. Available from <http://www.pcisig.com>.

PCI Power Management specification. Available from <http://www.pcisig.com>.

PCI Single Root I/O Virtualization, revision 1.1. Available from  
[http://www.pcisig.com/specifications/iov/single\\_root/](http://www.pcisig.com/specifications/iov/single_root/).

PCI Firmware 3.0 specification. Available from <http://www.pcisig.com>.

UEFI 2.3.1 specification. Available from <http://www.uefi.org>.

Trusted Computing Group Storage Architecture Core specification. Available from  
<http://www.trustedcomputinggroup.org>.

JEDEC JESD218: Solid State Drive (SSD) Requirements and Endurance Test Method standard. Available from <http://www.jedec.org>.

## Disposition log

7/11/2012	Erratum captured.
7/18/2012	Removed chapter 10, added references to UEFI to section 1.10.
8/28/2012	Erratum ratified.

*Technical input submitted to the NVMHCI Workgroup is subject to the terms of the NVMHCI Contributor's agreement.*